Composite Bayesian inference

Alexis Roche*
CHUV / Siemens Healthcare / EPFL
CH-1015 Lausanne, Switzerland

arXiv:1512.07678v3 [stat.CO] 3 May 2016

This paper revisits the concept of composite likelihood from the perspective of probabilistic inference, and proposes a generalization called *super composite likelihood* for sharper inference in multiclass problems. It is argued that, beside providing a new interpretation and a general justification of naïve Bayes procedures, super composite likelihood yields a much wider class of discriminative models suitable for unsupervised and weakly supervised learning.

1. Introduction

Conventional Bayesian inference and other likelihood-based paradigms rest upon the existence of a statistical data-generating model that is both experimentally plausible and computationally tractable. Because this is challenging when the data is inherently complex, common practice to implement feasible inference algorithms is to use deliberately misspecified data-generating models [27, 25] such as in *naïve Bayes* [18] or minimum description length [12], or to resort to highly supervised discriminative modeling approaches [18], not to mention ad-hoc methodology.

Composite likelihood (see [23] and the references therein) is a middle-way approach that extends the familiar notion of likelihood without requiring a full data-generating model. The key idea is to model an arbitrary set of low-dimensional features separately and then combine them, instead of modeling the data distribution as a whole. This may also be viewed as a *divide & conquer* method of approximating the actual likelihood. While maximum composite likelihood does not inherit the general property of maximum likelihood to yield asymptotically minimum-variance estimators, it may offer an excellent trade-off between computational and statistical efficiency.

In this note, composite likelihood is interpreted as a probabilistic opinion pool [11, 9] of "agents" using different pieces of information, or clues, extracted from the data. Each agent acts as a local Bayesian statistician expressing an opinion in the form of a posterior distribution on the unknown parameters of interest, or hypotheses, given a specific clue. Composite likelihood can therefore be associated with a probability distribution on hypotheses, hence extending Bayesian analysis to problems where the proper likelihood function is intractable.

I further justify a generalization of composite likelihood called *super composite likelihood* whereby clues can be weighted differently depending on hypotheses to better deal with multiclass inference problems. This more general concept also encompasses another likelihood approximation strategy known as *PDF projection* [2, 16, 3].

*alexis.roche@{centraliens.net,gmail.com,epfl.ch}

This paper is intended to serve as supporting material for a companion paper (in preparation), where the super composite likelihood framework is applied in computer vision to develop a probabilistic theory of image registration [24].

2. Composite likelihood as opinion pooling

Let Y an observable multivariate random variable with sampling distribution $p(y|\theta)$ conditional on some unobserved parameter of interest, $\theta \in \Theta$, where Θ is a known set. Given an experimental outcome y, the likelihood is the sampling distribution evaluated at y, seen as a function of θ:

$$L(\theta) = p(y|\theta).$$

For a high-dimensional Y, this expression may be intractable if a plausible generative model is lacking, or involves nuisance parameters that are cumbersome to integrate out. A natural workaround known as *data reduction* is then to extract some lower-dimensional feature $z = f(y)$, where f is a many-to-one mapping, and consider the potentially more convenient likelihood function:

$$\ell(\theta) = p(z|\theta).$$

Substituting $L(\theta)$ with $\ell(\theta)$ boils down to restricting the sample space, thereby "delegating" statistical inference to an "agent" provided with partial information. While it is valid for such an agent observing z only to consider $\ell(\theta)$ as the likelihood function of the problem, the drawback is that $\ell(\theta)$ might be poorly informative about θ due to the information loss incurred by data reduction. To make the trick statistically more efficient, we may extract several features, $z_i = f_i(y)$ for $i = 1, 2, \ldots, n$, and try to aggregate the likelihood functions $\ell_i(\theta) = p(z_i|\theta)$ that they elicit.

This leads to a classical problem of combining probabilistic opinions from possibly redundant agents [22, 11, 9, 1]. Genest *et al* [10] showed, in particular, that the only pooling operator that does not explicitly depend on θ and preserves *external Bayesianity* is the generalized logarithmic opinion pool:

$$p_\star(\theta) = \frac{1}{Z}\pi(\theta)\prod_{i=1}^{n}\ell_i(\theta)^{w_i}, \qquad \text{with} \quad Z = \int \pi(\theta)\prod_{i=1}^{n}\ell_i(\theta)^{w_i}d\theta, \tag{1}$$

where $\pi(\theta)$ is some reference distribution or prior, and w_i are arbitrary positive[1] weights that sum up to one,

$$\sum_{i=1}^{n} w_i = 1.$$

External Bayesianity essentially means that it should not matter whether a prior on θ is incorporated before or after pooling opinions, provided that all agents agree on the same prior. Importantly, the log-linear pool does not assume mutual feature independence, as would the same factorized form as (1) with all unitary weights, $w_1 = \ldots = w_n = 1$. Instead, redundancy between agents is assumed by default, and is effectively encoded by the unit sum constraint on weights[2].

[1] Negative weights can be chosen only if the parameter set Θ is finite [10].

[2] Nevertheless, features which are *known* to be mutually independent can be merged into a single feature. This results in increasing their weights in the log-linear pool.

Strikingly, (1) reduces to an analogous of Bayes rule: $p_\star(\theta) \propto \pi(\theta)L_c(\theta, \mathbf{w})$, where $\mathbf{w} = (w_1, w_2, \ldots, w_n)^\top$ denotes the vector of weights, and the quantity:

$$L_c(\theta, \mathbf{w}) \equiv \prod_{i=1}^{n} \ell_i(\theta)^{w_i} \tag{2}$$

plays exactly the same role as a traditional likelihood function. $L_c(\theta, \mathbf{w})$ shares a convenient factorized form with the likelihood derived under mutual feature independence, sometimes called *naïve Bayes* likelihood in the literature [18]. The key difference is that the single-feature likelihoods are scaled by positive weights w_i smaller than one, hence producing a flatter posterior distribution. In comparison with the proper likelihood (not assuming feature independence), the clear computational advantage is that we only need to evaluate the marginal feature distributions, rather than the joint distribution of all features.

We recognize in (2) a general expression known as a *marginal composite likelihood* [23], although it is derived here under the restriction that the weights sum up to one (as already motivated in [26] by a different argument using the maximum entropy principle). Owing to the opinion pooling interpretation, this simple constraint justifies plugging composite likelihood into Bayes rule. Previous attempts at Bayesian composite likelihood include tuning a constant weight so as to best adjust the pseudo posterior variance matrix to the asymptotic variance matrix of the maximum composite likelihood estimator [19], or performing a close-in-spirit curvature adjustment [21]. Such approaches reconcile the frequentist and Bayesian notions of uncertainty to some extent, but are not externally Bayesian since they do not warrant unit sum weights. When we refer to composite likelihood in the sequel, we assume unit sum weights.

3. Composite likelihood as message approximation

Composite likelihood may also be understood as a means to approximate the "true" likelihood function[3] or, in the language of graphical models, the *message* that the data sends to the latent variable θ. Several "clues" are sent to θ via different data features, and then integrated assuming *non-coalescent* emitting sources, meaning that statistical dependences between clues are treated as unknown, although not ignored. The whole idea is depicted by the factor graph[4] in Figure 1(b), where the latent variable is connected to multiple factors involving single clues (rather than a single factor involving multiple clues), thereby enabling efficient computations.

This approximation scheme happens to be optimal in an information-theoretic sense. As noted in [9], the log-linear pool minimizes the average Kullback-Leibler (KL) divergence to the probabilistic opinions:

$$p_\star = \arg\min_p \sum_{i=1}^{n} w_i D(p \| p_i), \tag{3}$$

where $p_i(\theta) \propto \pi(\theta)\ell_i(\theta)$ is the posterior distribution of agent i. Note that, even though the weights are arbitrary in (3), the optimal solution normalizes them to unit sum. Composite likelihood thus arises as the best possible consensus among agents according to a natural criterion, in addition to satisfying the external Bayesianity axiom.

If one interprets the average divergence from agent opinions (3) as a proxy for the divergence from "the truth", $p(\theta|y)$, composite likelihood can be seen as a variational approximation to

[3] By *true likelihood*, we mean the likelihood corresponding to a specified yet intractable generative model. In practice, such a model is obviously not required.

[4] See [6] for a didactic introduction to factor graphs.

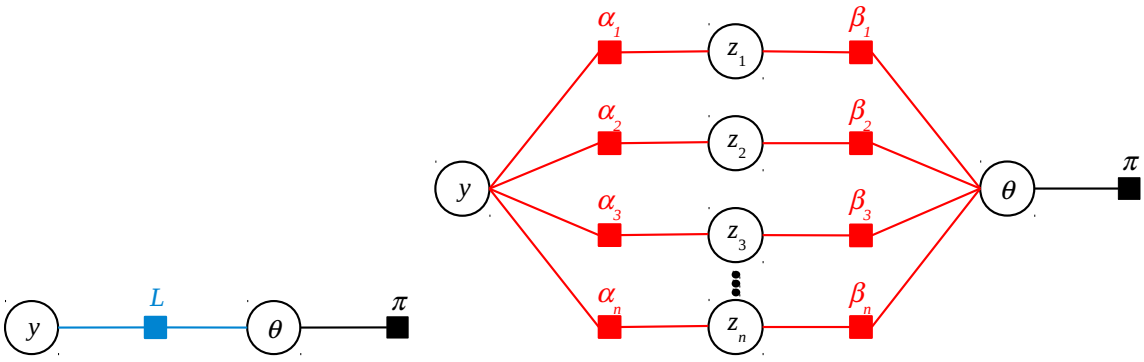

(a) Generative model (b) Composite likelihood model

Figure 1: Factor graph representations of a generative model (a) and its approximation using composite likelihood (b). In (a), the factor connecting the data y and the latent variable θ is the true likelihood, $L(y, \theta) = p(y|\theta)$. In (b), factors connecting y and the clues $z_1, z_2, \ldots$ represent feature extractions, $\alpha_i(y, z_i) = \delta[y - f_i(z_i)]$, while factors connecting the clues and θ are scaled single-feature likelihoods, $\beta_i(z_i, \theta) = p(z_i|\theta)^{w_i}$. In both graphs, the factor shown in black is the prior $\pi(\theta)$.

the true likelihood. This corresponds to the intuitive notion that a consensus among sufficiently many experts should yield a reasonable guess. An essential difference with usual approximate inference methods [6, 17] is that the likelihood function does not need to be computable owing to the use of a proxy. However, as we shall see in Section 5, knowledge of the feature sampling distributions can be further exploited to optimize the likelihood approximation with respect to the composite likelihood weights.

4. Super composite likelihood

Due to the distribution of weights between clues, a drawback of composite likelihood is that it is prone to *information overload* in the sense that it tends to "flatten out" when too many clues are included as relevant clues then get downweighted. If one decides not to merge clues for computational reasons (this would require handling their joint distribution), one could hope to mitigate information overload by assigning strong weights only to those clues that are believed to be "most informative".

However, when chosen for computational simplicity, clues may not only convey limited information at individual level: their informativeness may also be very much hypothesis-dependent. Consider, for instance, diagnosing a disease from a routine medical checkup. Body temperature may point to a bacterial infection by comparison with normality, but would not help detecting a non-infectious cardiovascular disease – and conversely for, say, blood pressure.

This motivates a more general setting where clues can be weighted differently depending on hypotheses. To avoid unnecessary technicalities, we will assume from now on a finite set of hypotheses, $\Theta = \{\theta_0, \theta_1, \ldots, \theta_m\}$, where one particular hypothesis, θ_0, is given the special status of reference, or "null" hypothesis.

We start by introducing auxiliary binary variables t_j, for $j = 1, \ldots, m$, defined by truncation

of θ:

$$t_j = \begin{cases} 1 & \text{if} \quad \theta = \theta_j \\ 0 & \text{if} \quad \theta = \theta_0 \end{cases} . \tag{4}$$

One can think of each t_j as an indicator light that flashes green or red whenever θ is in one of the particular two states θ_j or θ_0, and does not respond otherwise. The collection of all t_j's may be thought of as a population code [14] for θ.

The key idea is as follows: instead of approximating the message sent from the data to θ in one piece as in Section 3, we may approximate each of the simpler messages sent to the t_j's. To that end, we construct a factor graph using the population code, depicted in Figure 2, which is equivalent to the graph in Figure 1(a) for the purpose of computing the posterior distribution $p(\theta|y)$. This graph involves multiple factors representing the "truncated" likelihood functions,

$$L_j(t_j) = p(y|t_j) = \begin{cases} p(y|\theta_j) & \text{if} \quad t_j = 1 \\ p(y|\theta_0) & \text{if} \quad t_j = 0 \end{cases} , \qquad j = 1, \ldots, m,$$

as opposed to a single factor for the full likelihood $L(\theta)$. In addition, there are factors γ_j to synthesize the different messages sent to the binary variables $t_1, \ldots, t_m$ into one sent to θ:

$$\gamma_j(t_j, \theta) = \begin{cases} \delta_{1t_j} & \text{if} \quad \theta = \theta_j \\ \delta_{0t_j} & \text{if} \quad \theta \neq \theta_j \end{cases} , \tag{5}$$

where δ denotes the Kronecker delta.

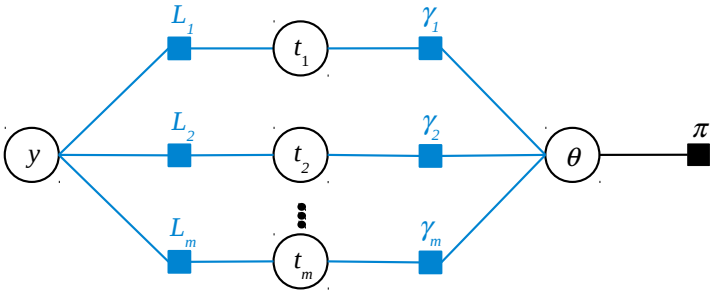

Figure 2: Alternative graph using population coding for the computation of $p(\theta|y)$. See text for details.

The unnormalized joint distribution $p_2(y, \theta)$ represented by the graph in Figure 2 reads:

$$\begin{aligned} p_2(y,\theta) &= \pi(\theta) \sum_{t_1=0}^{1} \cdots \sum_{t_m=0}^{1} \prod_{j=1}^{m} L_j(t_j)\gamma_j(t_j,\theta) \\ &= \pi(\theta) \prod_{j=1}^{m} \left\{ \mathbf{1}_{\{\theta_j\}}(\theta) p(y|\theta_j) + [1 - \mathbf{1}_{\{\theta_j\}}(\theta)] p(y|\theta_0) \right\} \\ &= \pi(\theta) p(y|\theta) p(y|\theta_0)^{m-1}, \end{aligned}$$

clearly yielding the same posterior $p_2(\theta|y) = p(\theta|y)$ as the graph in Figure 1(a), although generally not the same generative model: $p_2(y|\theta) \propto p(y|\theta)p(y|\theta_0)^{m-1}$ hence $p_2(y|\theta) \neq p(y|\theta)$ unless there are two hypotheses only ($m = 1$), or the null distribution $p(y|\theta_0)$ is uniform. While the inconsistency between generative models is irrelevant to inference, the alternative

graph enables a more flexible likelihood approximation scheme, where each factor $L_j(t_j)$ can be substituted with a specific composite likelihood $L_{cj}(t_j, \mathbf{w}_j)$ using own pre-determined weights $\mathbf{w}_j = (w_{1j}, w_{2j}, \ldots, w_{nj})^\top$ depending on j:

$$L_{cj}(t_j, \mathbf{w}_j) = \prod_{i=1}^{n} p(z_i|t_j)^{w_{ij}} = \prod_{i=1}^{n} \beta_{ij}(z_i, t_j),$$

with:

$$\beta_{ij}(z_i, t_j) = \begin{cases} p(z_i|\theta_j)^{w_{ij}} = \ell_i(\theta_j)^{w_{ij}} & \text{if} \quad t_j = 1 \\ p(z_i|\theta_0)^{w_{ij}} = \ell_i(\theta_0)^{w_{ij}} & \text{if} \quad t_j = 0 \end{cases}, \tag{6}$$

and:

$$\forall j \in \{1, \ldots, m\}, \qquad \sum_{i=1}^{n} w_{ij} = 1.$$

This corresponds to replacing each factor L_j in Figure 2 with a subgraph of same structure as the subgraph highlighted in red in Figure 1(b), resulting in the further modified factor graph shown in Figure 3. Intuitively, hypothesis-dependent weights make it possible to emphasize the clues that are relevant to each particular hypothesis comparison θ_j *vs.* θ_0, leading to potentially better approximations to the true odds $p(\theta_j|y)/p(\theta_0|y)$ than using constant weights.

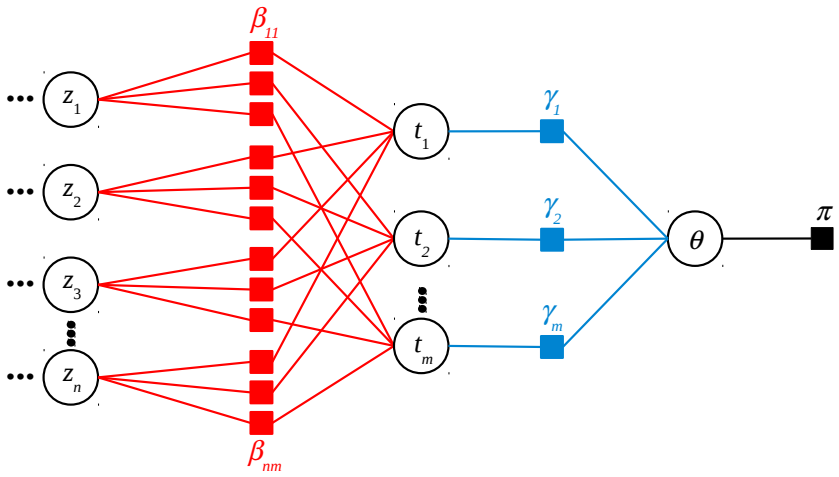

Figure 3: Super composite likelihood factor graph. See respectively (6) and (5) for the expression of the factors β_{ij} and γ_j.

We shall point out that the subgraph connecting $\mathbf{z} = (z_1, z_2, \ldots, z_n)$ and $\mathbf{t} = (t_1, t_2, \ldots, t_m)$ in Figure 3 (shown in red) is equivalent to an undirected bipartite graph, a property shared with restricted Boltzmann machines (RBM) [8, 13, 15]. This means that the variables in one layer are independent conditionally on the other layer:

$$p_3(\mathbf{t}|\mathbf{z}) = \prod_{j=1}^{m} p_3(t_j|\mathbf{z}), \qquad p_3(\mathbf{z}|\mathbf{t}) = \prod_{i=1}^{n} p_3(z_i|\mathbf{t}), \tag{7}$$

with, in this case,

$$p_3(t_j|\mathbf{z}) \propto L_{cj}(t_j, \mathbf{w}_j), \qquad p_3(z_i|\mathbf{t}) = p_3(z_i|\theta) \propto p(z_i|\theta)^{W_i(\theta)},$$

where $W_i(\theta)$ is defined by $W_i(\theta_j) = w_{ij}$ if $j \in \{1, m\}$ and $W_i(\theta_0) = \sum_{j=1}^m w_{ij}$.

An essential difference with RBM, however, is that the generative distribution $p_3(\mathbf{z}|\mathbf{t})$ is not intended to be used for model training. Instead, our core assumption is that the marginal feature distributions are learned "before" constructing the graph, a step that does not rely on assuming feature independence conditionally on $\mathbf{t}$ or, equivalently, θ. Therefore, despite satisfying the bipartite condition (7), the joint distribution $p_3(\mathbf{t}, \mathbf{z})$ represented in Figure 1(b) is asymmetrical in use, and serves the only purpose of defining a posterior distribution $p_3(\theta|y)$ via $p_3(\mathbf{t}|\mathbf{z})$.

Marginalizing out $\mathbf{t}$, we find:

$$\begin{aligned} p_3(\theta|y) &\propto \pi(\theta) \sum_{\mathbf{t}\in\{0,1\}^m} \prod_{j=1}^m p_3(\mathbf{t}|\mathbf{z})\gamma_j(t_j, \theta) \\ &= \pi(\theta) \sum_{t_1=0}^1 \dots \sum_{t_m=0}^1 \prod_{j=1}^m L_{cj}(t_j, \mathbf{w}_j)\gamma_j(t_j, \theta) \\ &= \pi(\theta) \prod_{i=1}^n \prod_{j=1}^m \ell_i \left(\mathbf{1}_{\{\theta_j\}}(\theta)\theta + [1 - \mathbf{1}_{\{\theta_j\}}(\theta)]\theta_0\right)^{w_{ij}} \\ &= \pi(\theta) \prod_{i=1}^n \ell_i(\theta_0)^{\sum_{j=1}^m w_{ij}} \prod_{i=1}^n \left[\frac{\ell_i(\theta)}{\ell_i(\theta_0)}\right]^{w_{ij}}, \end{aligned}$$

therefore:

$$p_3(\theta|y) = K \frac{\pi(\theta)}{\pi(\theta_0)} \prod_{i=1}^n \left[\frac{\ell_i(\theta)}{\ell_i(\theta_0)}\right]^{w_i(\theta)}, \tag{8}$$

where K is a normalizing factor (dependent on y) and the functions $w_i(\theta)$ are defined by $w_i(\theta_j) = w_{ij}$ for $j = 1, \dots, m$ and $w_i(\theta_0) = 0$ conventionally.

As an opinion pooling rule, (8) is more general than the log-linear pool (1) as it explicitly depends on θ through the changing weights. Moreover, since the weights sum up to one for each $\theta \neq \theta_0$, we have the equivalent expression:

$$p_3(\theta|y) = K \prod_{i=1}^n \left[\frac{\pi(\theta)\ell_i(\theta)}{\pi(\theta_0)\ell_i(\theta_0)}\right]^{w_i(\theta)},$$

showing that the pool does not change depending on whether the prior is incorporated before or after combining experts. In other words, (8) is, again, externally Bayesian.

As in Section 2, we recognize a Bayes rule-type expression: $p_3(\theta|y) \propto \pi(\theta)\mathcal{L}_c(\theta, \mathbf{W})$, with:

$$\mathcal{L}_c(\theta, \mathbf{W}) \equiv \prod_{i=1}^n \left[\frac{\ell_i(\theta)}{\ell_i(\theta_0)}\right]^{w_i(\theta)}, \tag{9}$$

where $\mathbf{W}$ denotes the $n \times m$ weight matrix with general element w_{ij}. We call (9) a super composite likelihood (SCL). Note that the SCL evaluated at a particular hypothesis θ_j boils down to a composite likelihood ratio,

$$\mathcal{L}_c(\theta_j, \mathbf{W}) = \frac{L_c(\theta_j, \mathbf{w}_j)}{L_c(\theta_0, , \mathbf{w}_j)}. \tag{10}$$

To see that SCL (9) is indeed a generalization of composite likelihood, assume that $\mathbf{w}_j = \mathbf{w}$ is the same for all hypotheses. This implies that (10) simplifies to $\mathcal{L}_c(\theta, \mathbf{W}) = L_c(\theta, \mathbf{w})/L_c(\theta_0, \mathbf{w})$,

the denominator of which is independent from θ and can therefore be safely ignored for inference about θ. In this special case, SCL is equivalent to standard composite likelihood regardless of the chosen reference θ_0.

Nevertheless, θ_0 plays a crucial normalization role whenever the columns of $\mathbf{W}$ are different to enable complex, possibly sparse, weighting patterns. Consider, for instance, the case where each hypothesis gets evidence from a single clue, so that $\mathbf{W}$ has a single 1 in each column and all other elements 0. The SCL (9) then reduces to a likelihood ratio involving a single, hypothesis-specific clue $z_{\iota(j)}$ determined by some mapping $\iota : \{1, \ldots, m\} \to \{1, \ldots, n\}$:

$$\mathcal{L}_c(\theta_j, \mathbf{W}) = \frac{\ell_{\iota(j)}(\theta_j)}{\ell_{\iota(j)}(\theta_0)} = \frac{p(z_{\iota(j)}|\theta_j)}{p(z_{\iota(j)}|\theta_0)}. \tag{11}$$

This expression was already used in [2, 16, 3] motivated by a different argument than here, namely the *PDF projection theorem*, which states the full data-generating model maximizing relative entropy (with respect to a chosen reference distribution) under knowledge of the feature sampling distributions. As discussed in [2], (11) closely approximates the true likelihood if each $z_{\iota(j)}$ is a near-sufficient statistic for θ_j *vs.* θ_0, a condition that may be difficult to meet if the choice of clues is driven by computational efficiency. SCL provides an alternative interpretation of PDF projection, while also extending it to multiple clues with unknown statistical dependences.

5. Weight optimization

The general advantage of SCL over standard composite likelihood is to define a broader class of pseudo-likelihood functions, hence with the potential to better approximate the true likelihood for a suitable choice of the weight matrix $\mathbf{W}$. On the other hand, standard composite likelihood satisfies good asymptotic properties for any choice of unit sum positive weights, as shown in Appendix B, but such properties do not extend in full generality to SCL: a poor likelihood approximation is to be expected if $\mathbf{W}$ is chosen at random, hence the importance of fine-tuning super composite weights.

To this end, a conventional machine learning strategy may be used. Assume (for now) that a number of examples of inputs and responses (y^k, θ^k), for $k = 1, \ldots, N$, are sampled independently. Given a tentative weight matrix $\mathbf{W}$, examples elicit different SCL functions $\mathcal{L}_{ck}(\theta, \mathbf{W})$ as y^k varies across examples. By analogy with classical maximum likelihood model selection, we may consider rating the model ability to predict examples depending on $\mathbf{W}$ via the sample average logarithm of the SCL:

$$\begin{aligned} \hat{\mathcal{U}}(\mathbf{W}) &= \frac{1}{N}\sum_{k=1}^{N} \log \mathcal{L}_{ck}(\theta^k, \mathbf{W}) \\ &= \frac{1}{N}\sum_{k=1}^{N}\sum_{i=1}^{n} w_i(\theta^k) \log \frac{p(z_i^k|\theta^k)}{p(z_i^k|\theta_0)}, \end{aligned}$$

which essentially averages feature-based log-likelihood ratios relative to the chosen reference hypothesis θ_0.

This utility measure turns out to be particularly appealing to tune SCL weights, for two reasons. First, if examples are drawn from the true yet incompletely known joint distribution $p(y, \theta) = \pi(\theta)p(y|\theta)$, then, under the condition of existence, $\hat{\mathcal{U}}(\mathbf{W})$ converges in the limit of

many examples to:

$$\begin{aligned}\mathcal{U}(\mathbf{W}) &= E\left[\log \mathcal{L}_c(\theta, \mathbf{W})\right] \\ &= \sum_{j=1}^{m} \pi(\theta_j) \sum_{i=1}^{n} w_{ij} \underbrace{\int p(z_i|\theta_j) \log \frac{p(z_i|\theta_j)}{p(z_i|\theta_0)} dz_i}_{u_{ij}}, \end{aligned} \tag{12}$$

where u_{ij} denotes the KL divergence of $p(z_i|\theta_0)$ from $p(z_i|\theta_j)$, and can be interpreted as the utility of clue i regarding hypothesis j. Since the u_{ij}'s are fully determined by the known feature-generating distributions $p(z_i|\theta)$, (12) can be evaluated without further knowledge about $p(y, \theta)$. Consequently, we do not need to constitute an actual training dataset!

Second, let $\mathcal{U}_\star$ denote the expected utility associated with the true distribution $p(y, \theta)$, considered as a SCL built from a single clue, the full data y. As is easy to check, $\mathcal{U}_\star$ is the conditional Kullback-Leibler divergence of $p(y|\theta_0)$ from $p(y|\theta)$:

$$\mathcal{U}_\star = \sum_{j=1}^{m} \pi(\theta_j) \int p(y|\theta_j) \log \frac{p(y|\theta_j)}{p(y|\theta_0)} dy = D[p(y|\theta) \| p(y|\theta_0)].$$

Applying in (12) the data reduction inequality recalled in Appendix A, we obtain the intuitive result that $\mathcal{U}(\mathbf{W})$ is upper bounded by the expected utility achievable without data reduction:

$$\mathcal{U}(\mathbf{W}) \leq \mathcal{U}_\star.$$

This means that the full data-generating model maximizes the expected log-SCL over the whole space of SCL functions (*i.e.*, SCL functions built from arbitrary clues in arbitrary number and using arbitrary weights). Therefore, (12) qualifies as a measure of goodness of fit to the true likelihood.

Maximizing (12) with respect to $\mathbf{W}$ amounts to optimizing the fit over a subspace of SCL functions spanned by fixed user-specified features. Some constraints may be imposed on $\mathbf{W}$, such as assuming identical columns as in the case of standard composite likelihood, or forcing some elements to zero if some feature likelihood functions are only known on subsets of Θ. We here focus on the case where no constraints are imposed beside, of course, that the columns of $\mathbf{W}$ lie in the n-dimensional simplex. Clearly, the optimal weights are then determined independently for each hypothesis (and thus independently from the prior) as any positive weights satisfying, for all $j \in \{1, \ldots, n\}$:

$$w_{ij} = 0 \quad \text{iff} \quad i \notin \arg \max_{i' \in \{1,\ldots,n\}} u_{i'j}, \qquad \text{and} \quad \sum_{i=1}^{n} w_{ij} = 1, \tag{13}$$

yielding a sparsity-enforcing rule which assigns non-zero weights only to the clues that achieve maximal KL utility for a particular hypothesis. If there is a unique KL-maximal, or "most exhaustive" clue for each hypothesis, the resulting SCL function essentially boils down to the PDF projection method [2, 16, 3]. More generally, if several clues z_i maximize u_{ij} for some hypothesis θ_j, then optimal weights are not unique. In such case, it is preferable to assign equal weights to the winning clues in order to ensure that utility is not only large in average, but also stable across examples (*i.e.*, low in variance).

Also note that the weight optimality rule (13) implies that, for any $\theta_\star \in \Theta$,

$$E\left[\log \frac{L_c(\theta_\star, \mathbf{w})}{L_c(\theta_0, \mathbf{w})}\right] \leq E\left[\log \frac{L_c(\theta_\star, \mathbf{w}_\star)}{L_c(\theta_0, \mathbf{w}_\star)}\right],$$

where $\mathbf{w}_\star$ is the weight vector associated with $\theta_\star$, E stands for the expectation with respect to $p(y|\theta_\star)$, and $\mathbf{w}$ is any weight vector. Moreover, the following inequality is shown in Appendix B:

$$E\left[\log \frac{L_c(\theta, \mathbf{w})}{L_c(\theta_0, \mathbf{w})}\right] \leq E\left[\log \frac{L_c(\theta_\star, \mathbf{w})}{L_c(\theta_0, \mathbf{w})}\right],$$

for any hypothesis θ and weighting $\mathbf{w}$ and, in particular, for the weight vector associated with θ via (13). Therefore, provided that the weight matrix $\mathbf{W}$ verifies (13), we have that:

$$E\left[\log \mathcal{L}_c(\theta, \mathbf{W})\right] \leq E\left[\log \mathcal{L}_c(\theta_\star, \mathbf{W})\right],$$

meaning that the SCL is asymptotically consistent in the sense that the expectation of its logarithm is maximized by the "true" parameter value. This is a most awaited frequentist property which, as already pointed out, holds with any choice of weights for standard composite likelihood but not, in general, for SCL.

6. Further extensions

6.1. Conditional super composite likelihood

The SCL derivation rests upon the definition of feature-based likelihood as $\ell_i(\theta) = p(z_i|\theta)$. As a straightforward extension, $\ell_i(\theta)$ may be conditioned by an additional "independent" feature $z_i^c = f_i^c(y)$ considered as a predictor of the "dependent" feature, $z_i = f_i(y)$, yielding the more general form:

$$\ell_i(\theta) = p(z_i|\theta, z_i^c). \tag{14}$$

Conditioning may be useful if it is believed that z_i^c alone is little or not informative about θ, but can provide relevant information when considered jointly with z_i, as in the case of regression covariates, for instance. Standard composite likelihood (2) then amounts to *conditional composite likelihood* [23], a more general form of composite likelihood also including Besag's historical *pseudo-likelihood* [5], which was a major breakthrough in computer vision.

Likewise, the above derivation remains valid when "independent" features are used, and we can thus define a conditional version of SCL by plugging likelihood functions of the form (14) into (9).

6.2. Nuisance parameters

Most often, plausible feature-generating models involve unknown quantities of no direct interest. Such quantities may be estimated offline in a supervised learning phase if a suitable training dataset is available. However, if such "pre-training" is not feasible in practice, the feature-based likelihoods depend on an unknown parameter ψ, in addition to the parameter of interest θ: $\ell_i(\theta, \psi) = p(z_i|\theta, \psi)$, or $\ell_i(\theta, \psi) = p(z_i|z_i^c, \theta, \psi)$ if "independent" features are used. We then face two difficulties:

- *Parameter integration.* How to make the inference on θ independent from ψ assuming known SCL weights?
- *Weight optimization.* How to optimize the SCL weights under unknown ψ?

We address these two points in the sequel.

6.2.1. Parameter integration

When weighting clues independently from the hypotheses, a joint composite likelihood on both parameters may be derived:

$$L_c(\theta, \psi, \mathbf{w}) = \prod_i \ell_i(\theta, \psi)^{w_i},$$

and further integrated with respect to some prior on the nuisance parameter to yield a function of θ only, which we may call the *composite evidence*:

$$\bar{L}_c(\theta, \mathbf{w}) = \int \pi(\psi) L_c(\theta, \psi, \mathbf{w}) d\psi. \tag{15}$$

This corresponds to using the factor graph represented in Figure 4 to approximate the message from the data to θ, where the factors β_i are now defined by $\beta_i(z_i, \theta, \psi) = \ell_i(\theta, \psi)^{w_i}$. To compute the associated posterior $p_4(\theta|y)$, ψ is treated as an auxiliary variable to be integrated out, leading to $p_4(\theta|y) \propto \pi(\theta)\bar{L}_c(\theta, \mathbf{w})$.

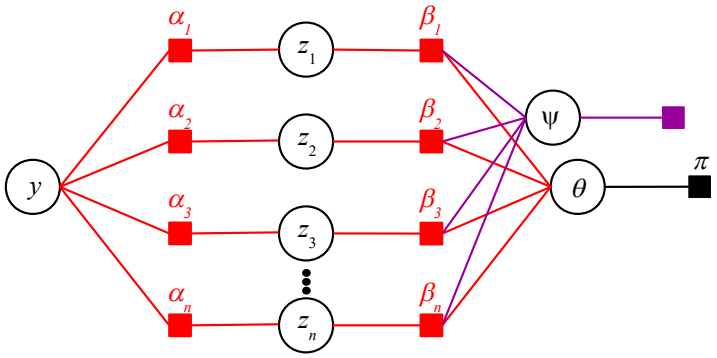

Figure 4: Extension of the composite likelihood factor graph in Figure 1(b) to account for nuisance parameters.

More generally, when weighting clues depending on hypotheses, we may use essentially the same idea as in Section 4, *i.e.*, replicate subgraphs such as the one depicted in red and magenta in Figure 4 in order to approximate the different messages sent from the data to each truncated variable t_j, as defined in (4). Note that, by replicating the same graph structure across variables t_j, the variable ψ is also replicated, so that the resulting factor graph in Figure 5 does not represent a joint distribution of the form $p(y, \theta, \psi)$, but rather one of the form $p(y, \theta, \psi_1, \ldots, \psi_m)$. This makes the approximation more flexible by taking advantage of the fact that a marginal distribution on ψ is not required.

The posterior distribution encoded by this graph (integrated with respect to the replicates $\psi_1, \ldots, \psi_m$) is easily found to be:

$$p_5(\theta|y) \propto \pi(\theta) \begin{cases} \dfrac{\bar{L}_c(\theta_j, \mathbf{w}_j)}{\bar{L}_c(\theta_0, \mathbf{w}_j)} & \text{if} \quad \theta = \theta_j \quad \text{with} \quad j \in \{1, \ldots, m\} \\ 1 & \text{if} \quad \theta = \theta_0 \end{cases},$$

where $\bar{L}_c(., \mathbf{w}_j)$ is the above-defined composite evidence function (15) associated with the weight

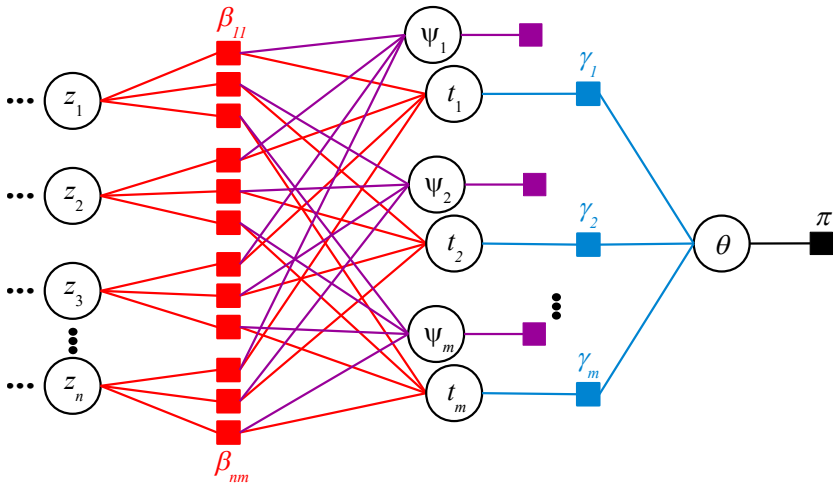

Figure 5: Extension of the super composite likelihood factor graph in Figure 3 to account for nuisance parameters. Here, the factors $\beta_{ij}(z_i, t_j, \psi_j)$ are defined by $\beta_{ij}(z_i, 1, \psi_j) = \ell_i(\theta_j, \psi_j)^{w_{ij}}$ and $\beta_{ij}(z_i, 0, \psi_j) = \ell_i(\theta_0, \psi_j)^{w_{ij}}$.

vector $\mathbf{w}_j$. This justifies defining the *super composite evidence* as:

$$\bar{\mathcal{L}}_c(\theta_j, \mathbf{W}) = \begin{cases} \dfrac{\bar{L}_c(\theta_j, \mathbf{w}_j)}{\bar{L}_c(\theta_0, \mathbf{w}_j)} & \text{if} \quad j \in \{1, \ldots, m\} \\ 1 & \text{if} \quad j = 0 \end{cases} . \tag{16}$$

Obvious from (16) is that the super composite evidence conserves all odds θ_j *vs.* θ_0 integrated within their respective associated subgraphs, very much like its nuisance parameter-free version (10). The conservation of odds implies that evaluating the SCL at a particular θ is akin to computing a scaled Bayes factor, which may sometimes be efficiently approximated via a maximum likelihood ratio statistic.

6.2.2. Weight optimization

Consistently with the idea of nuisance parameter replication (see Figure 4), we may employ a nested strategy to optimize the SCL weights in presence of nuisance parameters: for each $j \in \{1, \ldots, m\}$, apply the weight optimization of Section 5 to the subgraph model connecting all clues z_i with the particular pair (t_j, ψ_j). This leads to a set of linear programming problems: for each j, find the weight vector $\mathbf{w}_j$ on the n-dimensional simplex that maximizes:

$$\bar{\mathcal{U}}_j(\mathbf{w}_j) = E\left[\log \frac{L_c(\theta_j, \psi_j, \mathbf{w}_j)}{L_c(\theta_0, \psi_j, \mathbf{w}_j)}\right],$$

where the expectation is taken with respect to $(y, \psi_j) \sim p(y|\theta_j, \psi)\pi(\psi)$, a computation that only requires knowledge of the feature sampling distributions, $p(z_i|\theta, \psi)$ for $i = 1, \ldots, n$, in addition to the prior $\pi(\psi)$.

In this way, every $\mathbf{w}_j$ is optimal for the posterior of (t_j, ψ_j) or, more precisely, yields the best approximation to the true joint likelihood $L(\theta, \psi)$ for θ restricted to the binary set $\{\theta_0, \theta_j\}$. The ensuing marginal likelihood approximations are therefore as tight as possible for all hypothesis comparisons with θ_0, as is the goal of SCL weight optimization.

7. Discussion

Composite likelihood is a relatively recent concept from computational statistics that has mainly been developed so far in a frequentist perspective as a surrogate for the maximum likelihood method. In this paper, we have shown (to our best knowledge, for the first time) a deep connection between composite likelihood and probabilistic opinion pooling, thereby establishing composite likelihood as a class of *discriminative models*[5] for statistical inference and machine learning. This connection is possible under the mild restriction that composite likelihood weights are chosen to sum up to one.

7.1. From naïve Bayes to super composite likelihood

In the probabilistic opinion pooling perspective, composite likelihood is essentially a reinterpretation and a generalization of the *naïve Bayes* paradigm that relaxes the associated "naïve" mutual feature independence assumption. In particular, when all features are given equal weight, the composite likelihood is the naïve Bayes likelihood raised to the power $1/n$, where n is the number of selected features, yielding for instance the same maximum a posteriori (MAP) estimator if a flat prior is used. However, beside providing a simple justification to the wide use of naïve Bayes MAP algorithms, composite likelihood using unit sum weights also entails a conservative rescaling of credibility sets derived from naïve Bayes, which gets more drastic as the number of features increases. This comes as a consequence of relaxing feature independence.

We further argued that this rescaling may, to some extent, be unduly conservative in multiclass problems if features are weighted uniformly over the space of unobserved labels, leading us to propose the more general concept of super composite likelihood (SCL). SCL essentially approximates likelihood ratios relative to a fixed reference hypothesis using locally weighted composite likelihood functions. Owing to weight adaptability, SCL describes a more general class of discriminative models than standard composite likelihood.

The idea can also be understood in terms of approximating the factor graph in Figure 2, which encodes the true but intractable posterior distribution, by another factor graph of the type depicted in Figure 3. This substitution results in breaking into pieces both the observed and unobserved variable spaces, and assembling a series of concise *messages* passing from the former to the latter. Note that this is a pretty unusual case of approximate inference in factor graphs where the factors to be approximated are unknown (or need not be known).

7.2. Super composite likelihood training

Any SCL model is fully determined by the marginal generative distributions of some prespecified features, which may rely on a moderate number of parameters if low-dimensional features are chosen. There are two approaches to deal with these parameters.

One is to estimate them beforehand by supervised learning, if an adequate training dataset is available. This could be compared with contrastive pre-training of RBMs [13, 8], which also optimizes parameters for generation of observable features. An important difference, however, is that contrastive pre-training is unsupervised. On the other hand, SCL pre-training relies on weaker assumptions as it does not assume conditional feature independence unlike RBMs. Once pre-training is complete, the SCL weights may be tuned by maximum SCL, as shown in Section 5: this second training stage is, again, supervised in essence, but does not require

[5] We here define a discriminative model as a model of which the parameters describe the conditional distribution of an unobserved variable of interest given an observable variable, as opposed to a *generative* model, which involves parameters encoding for the conditional distribution of an observable given an unobserved variable.

additional training examples since it is fully determined by the feature distributions learned in the pre-training step.

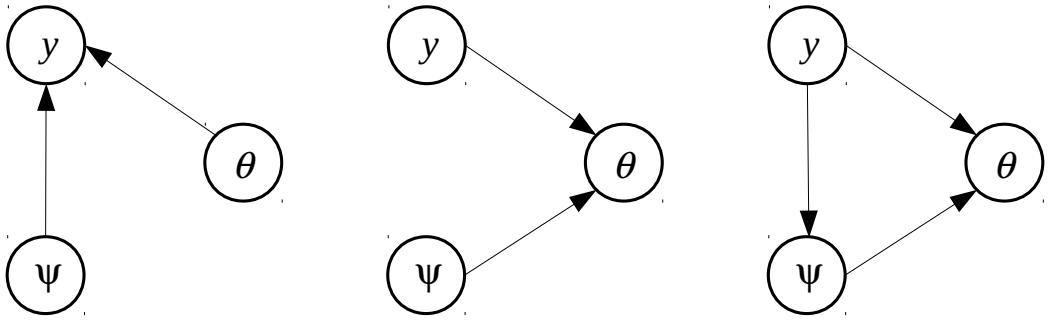

(a) Generative model

(b) Classical discriminative model

(c) Composite likelihood model

Figure 6: Belief networks representing, respectively: (a) a generative model, (b) a classical discriminative model, and (c) a composite likelihood model. Note the marginal independence between the data y and the nuisance parameter ψ in (b).

The other approach to deal with feature distribution parameters is to consider them as nuisance parameters, as proposed in Section 6.2, thereby avoiding pre-training. The method then becomes completely unsupervised. The weights can be tuned by maximizing the log-SCL integrated over the prior on nuisance parameters (see Section 6.2.2), which requires no training dataset whatsoever, and the nuisance parameters can be eliminated by substituting the SCL function with a *super composite evidence* (see Section 6.2.1). In this version, SCL essentially generalizes Bayesian integration.

The potential of SCL for unsupervised learning is perhaps suprising considering that it is a discriminative model. This stems from the fact that, unlike classical discriminative models, SCL exploits direct information conveyed by the data about both the parameters of interest and the nuisance parameters, reflecting the generative nature of the underlying feature distribution models. This essential difference is illustrated in Figure 6: the factor graph model underlying composite likelihood in its parametric version (see Figure 4) is not reducible to a belief network of the type represented in Figure 6(b), in which the data and the nuisance parameter are marginally independent. In contrast, the data is informative about the nuisance parameter in a composite likelihood model, as shown in Figure 6(c).

8. Conclusion

In summary, (super) composite likelihood has the potential to yield weakly supervised or unsupervised Bayesian-like inference procedures depending on the particular task at hand. This property reflects the encoding of statistical relationships between the data and *all* unknown parameters. Composite likelihood thus appears as a trade-off between generative models, which are optimal for unsupervised learning but possibly intractable, and traditional discriminative models (logistic regression, Gaussian processes [20], maximum entropy models [4], etc.), which are inherently supervised. Composite likelihood models are discriminative models assembled from atomic generative models and, from this point of view, may be considered as *semi-generative* models.

A. Data reduction inequality

A fundamental property of the Kullback-Leibler divergence is that it decreases under feature extraction, *i.e.*, application of a deterministic transformation. This comes as a consequence of the logarithmic sum inequality [7] (and also follows as a straightforward corollary of the PDF projection theorem [16, 3]).

The proof is elementary and can be sketched as follows. Let $Z = f(Y)$ some feature extracted from a variable $Y \sim p(y)$. Given an arbitrary distribution $\pi(y)$, let $\tilde{p}(z)$ and $\tilde{\pi}(z)$ the distributions induced on Z by $p(y)$ and $\pi(y)$, respectively. For any potential value z of Z, consider the level set: $\Gamma(z) = \{y, f(y) = z\}$. Under conditions of existence, we can partition the integral involved in the KL divergence $D(p\|\pi)$ using these level sets:

$$\int p(y) \log \frac{p(y)}{\pi(y)} dy = \int \left(\int_{\Gamma(z)} p(y) \log \frac{p(y)}{\pi(y)} dy \right) dz.$$

Applying the logarithmic sum inequality [7] to each integral over $\Gamma(z)$, we then readily get:

$$\int_{\Gamma(z)} p(y) \log \frac{p(y)}{\pi(y)} dy \geq \tilde{p}(z) \log \frac{\tilde{p}(z)}{\tilde{\pi}(z)},$$

owing to the fact that $\displaystyle\int_{\Gamma(z)} p(y)dy = \tilde{p}(z)$ by definition. Therefore,

$$\int \tilde{p}(z) \log \frac{\tilde{p}(z)}{\tilde{\pi}(z)} dz \leq \int p(y) \log \frac{p(y)}{\pi(y)} dy,$$

or, in a more compact way, $D(\tilde{p}\|\tilde{\pi}) \leq D(p\|\pi)$. The case of equality occurs if, and only if, $p(y)/\pi(y) = \tilde{p}(z)/\tilde{\pi}(z)$, in other words, if z is a sufficient statistic for the "alternative" hypothesis $H_1 : y \sim p(y)$ versus the "null" hypothesis $H_0 : y \sim \pi(y)$.

Since it is a well-known fact that $D(\tilde{p}\|\tilde{\pi}) \geq 0$ as any KL divergence, we conclude that:

$$0 \leq D(\tilde{p}\|\tilde{\pi}) \leq D(p\|\pi). \tag{17}$$

B. Basic asymptotic properties of composite likelihood

It follows from the double inequality (17) that, for any extracted feature z_i and for any hypothesis θ,

$$0 \leq E\left[\log \frac{p(z_i|\theta_\star)}{p(z_i|\theta)}\right] \leq E\left[\log \frac{p(y|\theta_\star)}{p(y|\theta)}\right],$$

where the expectation is taken with respect to the true distribution $p(y|\theta_\star)$. Using a weighted sum of such inequalities, we can bracket the expected variations of the logarithm of any standard composite likelihood function (assuming unit sum positive weights):

$$0 \leq E\left[\log \frac{L_c(\theta_\star, \mathbf{w})}{L_c(\theta, \mathbf{w})}\right] \leq E\left[\log \frac{L(\theta_\star)}{L(\theta)}\right]. \tag{18}$$

This implies two asymptotic properties of standard composite likelihood:

- *Consistency.* The expected log-composite likelihood is maximized by $\theta_\star$, the true value of θ.

- *Conservativeness.* The expected log-composite likelihood function is "flatter" than the true expected log-likelihood. In other words, composite likelihood ratios $L_c(\theta, \mathbf{w})/L_c(\theta_\star, \mathbf{w})$ relative to $\theta_\star$ tend to be larger than the corresponding true likelihood ratios.

These properties do not necessarily extend to the more general case of SCL. We may state a weaker consistency property by considering the upper envelope of expected log-composite likelihood ratios:

$$M(\theta) = \max_{\mathbf{w} \in \mathcal{W}} E\left[\log \frac{L_c(\theta, \mathbf{w})}{L_c(\theta_0, \mathbf{w})}\right],$$

where $\mathcal{W}$ is the n-dimensional simplex. Clearly, the expected logarithm of the SCL is upper bounded by $M(\theta)$:

$$E[\log \mathcal{L}_c(\theta, \mathbf{W})] \leq M(\theta).$$

Moreover, $M(\theta)$ is maximized by $\theta_\star$ since (18) implies that, for any (θ, θ_0),

$$E\left[\log \frac{L_c(\theta, \mathbf{w})}{L_c(\theta_0, \mathbf{w})}\right] \leq E\left[\log \frac{L_c(\theta_\star, \mathbf{w})}{L_c(\theta_0, \mathbf{w})}\right],$$

which remains true when taking the maximum over the weights in both sides:

$$M(\theta) = \max_{\mathbf{w} \in \mathcal{W}} E\left[\log \frac{L_c(\theta, \mathbf{w})}{L_c(\theta_0, \mathbf{w})}\right] \leq \max_{\mathbf{w} \in \mathcal{W}} E\left[\log \frac{L_c(\theta_\star, \mathbf{w})}{L_c(\theta_0, \mathbf{w})}\right] = M(\theta_\star).$$

Therefore, maximizing SCL corresponds to maximizing a lower bound on $M(\theta)$, which can be considered as an objective function since it is maximized by $\theta_\star$ (like the expected log-likelihood). Lower bound maximization guarantees a certain objective value, however not the maximum, hence asymptotic consistency may not hold.

References

[1] D. Allard, A. Comunian, and P. Renard. Probability aggregation methods in Geoscience. *Mathematical Geosciences*, 44(5):545–581, 2012.

[2] P. Baggenstoss. The PDF projection theorem and the class-specific method. *IEEE Transactions on Signal Processing*, 51(3):672–685, 2003.

[3] P. Baggenstoss. Maximum entropy PDF design using feature density constraints: applications in signal processing. *IEEE Transactions on Signal Processing*, 63(11):2815–2825, 2015.

[4] A. L. Berger, S. D. Pietra, and V. J. D. Pietra. A maximum entropy approach to natural language processing. *Computational Linguistics*, 22(1):39–71, 1996.

[5] J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B*, 36(2):192–236, 1974.

[6] C. Bishop. *Pattern Recognition and Machine Intelligence*. Springer, 2006.

[7] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., second edition, 2006.

[8] A. Fischer and C. Igel. Training restricted Boltzmann machines: an introduction. *Pattern Recognition*, 47(1):25–39, 2014.

[9] A. Garg, T. Jayram, S. Vaithyanathan, and H. Zhu. Generalized opinion pooling. In *8th International Symposium on Artificial Intelligence and Mathematics*, pages 79–86, 2004.

[10] C. Genest, K. J. McConway, and M. J. Schervish. Characterization of externally Bayesian pooling operators. *Annals of Statistics*, 14(2):487–501, 1986.

[11] C. Genest and J. V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–135, 1986.

[12] P. Grünwald. *The Minimum Description Length Principle*. MIT Press, 2007.

[13] G. E. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.

[14] D. C. Knill and A. Pouget. The Bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12):712–719, 2004.

[15] H. Larochelle and Y. Bengio. Classification using discriminative restricted Boltzmann machines. In *Proceedings of the 25th international conference on Machine learning*, pages 536–543. ACM, 2008.

[16] T. P. Minka. Exemplar-based likelihoods using the PDF projection theorem. Technical Report MSR-TR-2004-148, Microsoft Research Ltd., Cambridge, UK, 2004.

[17] T. P. Minka. Divergence measures and message passing. Technical Report MSR-TR-2005-173, Microsoft Research Ltd., Cambridge, UK, 2005.

[18] A. Y. Ng and M. I. Jordan. On discriminative vs. generative classifiers: a comparison of logistic regression and naive Bayes. In *Advances in Neural Information Processing Systems*, pages 841–848, 2001.

[19] F. Pauli, W. Racugno, and L. Ventura. Bayesian composite marginal likelihoods. *Statistica Sinica*, 21:149–164, 2011.

[20] C. Rasmussen and C. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.

[21] M. Ribatet, D. Cooley, and A. C. Davison. Bayesian inference from composite likelihoods, with an application to spatial extremes. *Statistica Sinica*, 22:813–845, 2012.

[22] A. Tarantola and B. Valette. Inverse problems = quest for information. *Journal of Geophysics*, 50(3):150–170, 1982.

[23] C. Varin, N. Reid, and D. Firth. An overview of composite likelihood methods. *Statistica Sinica*, 21(1):5–42, 2011.

[24] P. Viola and W. M. Wells. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24(2):137–154, 1997.

[25] S. G. Walker. Bayesian inference with misspecified models. *Journal of Statistical Planning and Inference*, 143(10):1621–1633, 2013.

[26] X. Wang and Y. Wu. Theoretical properties of composite likelihood. *Open Journal of Statistics*, 4:188–197, 2014.

[27] H. White. Maximum likelihood estimation of misspecified models. *Econometrica*, 50(1):1–25, 1982.